

Visual annotations and a supervised learning approach for evaluating and calibrating ChIP-seq peak detectors

Toby Dylan Hocking (toby.hocking@mail.mcgill.ca),
Patricia Goerner-Potvin, Andreanne Morin, Xiaojian Shao, and Guillaume Bourque
Department of Human Genetics, McGill University, Montréal, Canada.
McGill University and Génome Québec Innovation Center, Montréal, Canada.

September 23, 2014

Abstract

Many peak detection algorithms have been proposed for ChIP-seq data analysis, but it is not obvious which method and what parameters are optimal for any given data set. In contrast, peaks can easily be located by visual inspection of profile data on a genome browser. We thus propose a supervised machine learning approach to ChIP-seq data analysis, using annotated regions that encode an expert's qualitative judgments about which regions contain or do not contain peaks. The main idea is to manually annotate a small subset of the genome, and then learn a model that makes consistent predictions on the rest of the genome. We show how our method can be used to quantitatively calibrate and benchmark the performance of peak detection algorithms on specific data sets. We compare several peak detectors on 7 annotated region data sets, consisting of 2 histone marks, 4 expert annotators, and several different cell types. In these data the macs algorithm was best for a narrow peak histone profile (H3K4me3) while the hmcen.broad algorithm was best for a broad histone profile (H3K36me3). Our benchmark annotated region data sets can be downloaded from a public website, and there is an R package for computing the annotation error on GitHub.

Contents

1	Introduction and related work	3
1.1	Related work: benchmarking peak detectors	3
1.2	Related work: supervised, interactive analysis	4
2	Materials and Methods	4
2.1	Annotating samples	4
2.2	Annotation error function and peak detection problem	5
2.3	Calibrating a peak detection algorithm	6
2.4	Specific peak detection algorithms and parameters	6
2.5	Software availability	7
3	Results	7
3.1	Creating a benchmark data set of annotated regions	7
3.2	Train error for different mark types and annotators	9
3.3	Consistency of different annotators	10
3.4	Test on different regions of the same samples	10
3.5	Test on different samples and cell types	10
4	Discussion	11
4.1	Supervised versus unsupervised analysis	11
4.2	Time required for annotation	11
5	Conclusions	11
6	Acknowledgements	12
6.1	Funding	12
6.2	Conflict of interest statement.	12
A	Source code and data links	14
A.1	Original ChIP-seq data profiles	14
A.2	Annotated region database	14
A.3	R package with C code for computing the annotation error	14
A.4	L ^A T _E X and R source code for this article	14
B	Example of annotated region data file	14
C	Definition of the annotation error	15
C.1	Data definition	15
C.2	Weak annotation error	15
C.3	Strong annotation error	15
C.4	Total annotation error	16
C.5	ROC analysis	16
D	Details of peak calling algorithms	17
E	Supplementary Figures	18

1 Introduction and related work

Chromatin immunoprecipitation sequencing (ChIP-seq) is a genome-wide assay to profile histone modifications and transcription factor binding sites [Barski et al., 2007], with many experimental and computational steps [Bailey et al., 2013]. In this paper we propose a new method for the peak calling step. The goal of peak calling is to filter out background noise and accurately identify the locations of peaks in the genome.

There are two main lines of research into software tools that can help scientists find peaks in the genome. One class of software consists of peak detection algorithms, which are non-interactive command line programs that can be systematically run on all samples in a data set. An algorithm takes the aligned sequences as input, and returns precise locations of predicted peaks as output. Despite these advantages, peak detector software has one major drawback: model selection. There are many different algorithms that are specifically designed for peak detection in ChIP-seq data, each with many parameters to tune in each algorithm. Each algorithm and parameter combination will return a different set of predicted peaks. Given a specific ChIP-seq data set to analyze, how do you choose the best peak detection algorithm and its parameters?

The second class of software consists of graphical tools such as the UCSC genome browser [Kent et al., 2002]. To view ChIP-seq data on the UCSC genome browser, the ChIP-seq coverage must be saved to a bigWig file [Kent et al., 2010], which can be browsed as a line or bar plot to visually identify peaks. The main advantage of this approach to peak detection is that it is often easy to visually identify peaks and background noise in coverage plots of several ChIP-seq samples. There are two main disadvantages of this approach. First, precise peak start and end locations are not obvious on visual inspection. Second, no researcher has enough time to visually inspect and identify peaks across the whole genome.

In this article we propose a new machine learning approach for ChIP-seq data analysis that combines these two lines of research. The main idea is to manually annotate peaks in a small subset of the genome, and use those annotations to learn a peak detection model that makes consistent predictions on the rest of the genome. In particular, we propose to create annotated regions that encode an experienced scientist’s judgment about which regions contain or do not contain peaks (Figure 1). The annotated regions can then be used as a gold standard to calibrate model parameters and then evaluate peak detection algorithms on specific data sets.

1.1 Related work: benchmarking peak detectors

There are several algorithms for detecting peaks in ChIP-seq data [Zhang et al., 2008, Heinz et al., 2010, Ashoor et al., 2013, Song and Smith, 2011, Zang et al., 2009], and in this article we propose to benchmark their peak detection accuracy on specific data sets using manually annotated regions. Other methods for benchmarking ChIP-seq peak callers include using known binding sites [Chen and Zhang, 2010], low-throughput experiments [Osmanbeyoglu et al., 2012], and simulation studies [Szalkowski and Schmid, 2011]. Each of these benchmarking methods has its own strengths and weaknesses. For example, known binding sites are useful positive controls for transcription factor ChIP-seq, but are rarely known for histone marks. An unlimited amount of data can be generated using computational simulation studies, but these data may be arbitrarily different from real data sets of interest. Low-throughput experiments are always useful to confirm binding sites at specific genomic locations, but are not routinely done to accompany genome-wide ChIP-seq experiments.

The manual annotation method that we propose in this paper is much more widely applicable for benchmarking ChIP-seq peak detectors. The main weakness of our proposed method would be the time required for manual visual annotation of some genomic regions of the ChIP-seq data. However, we show in the results section that our method is useful even if there is only time to create a small annotated region database.

1.2 Related work: supervised, interactive analysis

Supervised machine learning methods have been applied to ChIP-seq without using interactive data visualization. For example, low-throughput experiments were used to train several supervised classification models to improve ChIP-seq peak calling [Osmanbeyoglu et al., 2012]. As another example, a supervised machine learning approach was used to define a regulatory vocabulary with genome-wide predictive power [Gorkin et al., 2012].

Other recent software tools focus on interactive visualization of ChIP-seq data, without using annotations and supervised machine learning approaches [Nielsen et al., 2012, Younesy et al., 2013]. They take several profiles and several genomic regions as input and allow the user to interactively adjust clustering model parameters. Both are similar to the approach used for Fluorescence-Activated Cell Sorting (FACS) data analysis, where the user manually specifies fluorescence thresholds for sorting cells.

Our interactive, visual approach to ChIP-seq data analysis is closely related to several other recently proposed software tools for supervised machine learning of biological data. For example, CellProfiler Analyst is an interactive system for semi-automatically labeling cell phenotypes in high-content cell microscopy screening assays [Jones et al., 2009].

To apply these visual methods to genomics data, we earlier proposed to benchmark breakpoint detection algorithms for DNA copy number analysis using a database of annotated regions that contain or do not contain breakpoints [Hocking et al., 2013]. We further proposed to use these annotated regions in SegAnnDB, a web site for supervised, interactive DNA copy number analysis [Hocking et al., 2014]. In the present paper we adapt this line of research for peak detection in ChIP-seq data.

2 Materials and Methods

2.1 Annotating samples

The first step of analysis is to create a database of annotated regions, by visually inspecting coverage plots of the ChIP-seq samples. It is advantageous to simultaneously inspect several samples of a single experiment type, to more easily identify common peaks and noise across multiple samples (Figure 1). It is useful to also simultaneously view supplementary data tracks such as genes, alignability/mappability, and related input/control samples. To ensure the creation of a gold standard annotated region database of high quality,

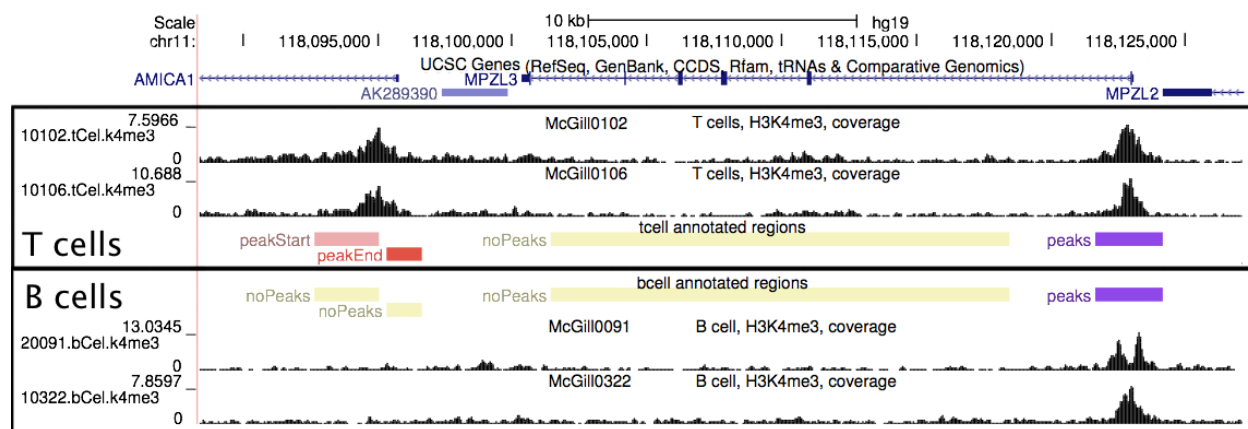


Figure 1: One genomic window containing 4 annotated regions for each of the 4 shown profiles. Visual inspection of ChIP-seq normalized coverage plots can be used to create annotated regions that encode where peaks should and should not be detected in these T cell and B cell samples. Exactly 1 peak start/end should be detected in each peakStart/peakEnd region. There should be no overlapping peaks in each noPeaks region, and at least 1 overlapping peak in each peaks region.

we only annotated regions which were very obvious, and avoided annotating any regions which were unclear. The boundaries of the regions can be made as large or small as necessary.

For each annotated region, we copied the genomic coordinates to a text file and noted the annotation (example file in supplementary materials). As shown in Table 1, we considered four types of annotations: peakStart, peakEnd, peaks, and noPeaks. Each noPeaks annotation is used to designate a region that definitely contains only background noise, and contains no peaks. When a peak start or end is visible, it can be annotated using a peakStart or peakEnd region. These regions should contain exactly one peak start or end, and do not necessarily need to occur in pairs. For example, if the peak start is clear and the peak end is unclear, then one should add a peakStart region, and not add any nearby peakEnd region. Finally, a peaks region means that there is at least one overlapping peak. These peaks regions are useful when the number of peaks, or the start and end locations are unclear. For example, on the right side of Figure 1 it is clear that there is at least one peak, but in profile McGill0091 there seems to be two peaks. So we created a peaks region, which means that either one or two peaks in that region is acceptable (but zero peaks is unacceptable).

We observed similar peaks across samples of a given cell type, and always assigned the same annotated regions to each of those samples (Figure 1). So that we could later review and verify the annotated regions, we grouped our annotation database into windows of nearby regions, such as the window shown in Figure 1. We made sure that each window contains at least one region with a peak and one region without peaks. Furthermore, we made sure that no region overlaps any other region on the same sample.

Finally, because each window contains the same annotated regions for samples of the same cell type, it is appropriate to use windows rather than regions as units of cross-validation. So in the computational experiments where we measure the peak detection test error, we train on several windows of annotated regions, and test on several other windows.

2.2 Annotation error function and peak detection problem

Assume we have n annotated training samples, all of the same ChIP-seq experiment type. For simplicity, and without loss of generality, let us consider just one chromosome with d base pairs. Let $\mathbf{x}_1 \in \mathbb{R}^d, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be the vectors of coverage across that chromosome. For example $d = 249,250,621$ is the number of base pairs on chr1, and $\mathbf{x}_i \in \mathbb{R}^d$ is the H3K4me3 coverage profile on chr1 for each sample $i \in \{1, \dots, n\}$.

We also have exactly four sets of annotated regions $\underline{R}_i, \bar{R}_i, R_i^+, R_i^-$ for each sample $i \in \{1, \dots, n\}$ (Table 1). Each region $\mathbf{r} \in R_i^+$ is an interval of base pairs. For example, Figure 2 shows a coverage profile that has one annotation of each type.




Figure 2: Annotated regions can be used to quantify the accuracy of a peak detection model. Peaks detected by five different thresholds of the HMCAN model are shown in blue for tcell sample McGill0102, mark H3K4me3, annotated by TDH. Models with too few peaks have false negatives (threshold ≥ 4), and models with too many peaks have false positives (threshold ≤ 1.5), so for these data we choose an intermediate threshold=2 that minimizes the number of incorrect annotations.

Table 1: Symbols and colors used to represent ChIP-seq coverage and annotated regions for n samples.

Data	Type	Color	Symbols
Coverage			$\mathbf{x}_1, \dots, \mathbf{x}_n$
Weak annotations	peaks	■	R_1^+, \dots, R_n^+
	noPeaks	■	R_1^-, \dots, R_n^-
Strong annotations	peakStart	■	$\underline{R}_1, \dots, \underline{R}_n$
	peakEnd	■	$\overline{R}_1, \dots, \overline{R}_n$

A peak detection function or peak caller $c : \mathbb{R}^d \rightarrow \{0, 1\}^d$ takes a coverage profile $\mathbf{x} \in \mathbb{R}^d$ as input, and returns a binary peak call prediction $\mathbf{y} = c(\mathbf{x}) \in \{0, 1\}^d$ (0 is background noise, 1 is a peak).

The goal is to learn how to call peaks $c(\mathbf{x}_i)$ which agree with the annotated regions $R_i^+, R_i^-, \underline{R}_i, \overline{R}_i$ for some test samples i . To quantify the error of the peak calls with respect to the annotation data, we define the annotation error as the sum of false positive (FP) and false negative (FN) regions:

$$E(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+, R_i^-) = \text{FP}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^-) + \text{FN}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+). \quad (1)$$

The principle used to compute the annotation error E is illustrated in Figure 2, and the precise mathematical definitions of FP and FN are given in the supplementary materials. In short, a false positive occurs when a peak detector predicts too many peaks in an annotated region, and a false negative occurs when there are not enough predicted peaks.

The supervised machine learning problem can be formalized as the following optimization problem. Find the peak caller c with minimal annotation error on a set of test samples:

$$\underset{c}{\text{minimize}} \sum_{i \in \text{test}} E[c(\mathbf{x}_i), \underline{R}_i, \overline{R}_i, R_i^+, R_i^-]. \quad (2)$$

In other words, the final goal is minimize the number of incorrect annotated regions on a test set of data. Note that “test” can be several different kinds of data sets, depending on how the peak detection model will be used (see Results section and Table 2).

2.3 Calibrating a peak detection algorithm

A standard unsupervised peak caller can be characterized as a function $c_\lambda : \mathbb{R}^d \rightarrow \{0, 1\}^d$, where the significance threshold $\lambda \in \mathbb{R}$ controls the number of peaks detected. In each peak detection algorithm, λ has a different, precise meaning that we specify in the supplementary materials. As shown in Figure 3, we select an optimal threshold λ by minimizing the annotation error on the set of n training samples

$$\hat{\lambda} = \arg \min_{\lambda} \sum_{i \in \{1, \dots, n\}} E[c_\lambda(\mathbf{x}_i), \underline{R}_i, \overline{R}_i, R_i^+, R_i^-]. \quad (3)$$

The training or model calibration procedure (3) consists of simply computing peak calls for several peak detection thresholds λ , and choosing whichever threshold $\hat{\lambda}$ minimizes the number of incorrect annotated regions. Or if there are no annotated regions available, we can simply use the default significance threshold $\tilde{\lambda}$ suggested by the author of each algorithm. The test error (2) can be used to evaluate the accuracy of the trained model $\hat{\lambda}$ and the default model $\tilde{\lambda}$.

2.4 Specific peak detection algorithms and parameters

We considered the following algorithms with free software implementations from the bioinformatics literature. We chose to compare several algorithms designed for peaks (such as H3K4me3) and several algorithms for

broad domains (such as H3K36me3). However, our analysis could be conducted with any peak detection algorithm, and we hope that authors of future algorithms will test them using annotated region databases.

MACS and HMCAN [Zhang et al., 2008, Ashoor et al., 2013] have parameters for peaks and broad domains, so we tried both settings. RSEG and SICER were designed for broad domains [Song and Smith, 2011, Zang et al., 2009]. The HOMER set of tools contains a findPeaks program which has been used to detect transcription factor binding sites and histone modifications [Heinz et al., 2010]. Details of each algorithm are given in the supplementary materials.

Each algorithm has several parameters that may affect peak detection accuracy. An exhaustive grid search over several parameters would be infeasible, since there is an exponential number of different parameter combinations. So for each algorithm we use only one parameter as the significance threshold λ , and held other parameters at default values. The precise meaning of the threshold λ for each algorithm is discussed in the supplementary materials.

2.5 Software availability

The annotation error (14) can be easily computed using the C code in the R package PeakError on GitHub: <https://github.com/tdhock/PeakError>

The `PeakError` function takes as input the annotated regions and peak predictions for one sample, and outputs the annotation error for each region.

3 Results

3.1 Creating a benchmark data set of annotated regions

We analyzed samples from the McGill Epigenomes portal (<http://epigenomesportal.ca>), which we visualized using the UCSC genome browser software [Kent et al., 2002]. We chose to annotate all samples available with data on H3K4me3 (representing profiles with narrow peaks) and all samples available with data on H3K36me3 (representing broadly enriched regions). In these data there are many samples of the same cell type, and the same peaks often occur in all samples of the same cell type (Figure 1). In total there were 37 samples with H3K4me3 data and 29 samples with H3K36me3 data, across 8 cell types (Table 3).

We constructed 7 annotated region databases (Table 3), using the method described in the “Annotating samples” section. Of the 4 different annotators, some were post-docs (TDH, XJ), and some were PhD students (AM, PGP). In total we created 12,826 annotated regions.




Figure 3: The annotated regions define an annotation error function that can be minimized (3) to choose an optimal parameter for a standard peak detector. We plot the percent error of the `hmcan.broad` model over 1743 annotated regions in the H3K36me3_AM_immune data set (other algorithms and data sets shown in Supplementary Figure S1).




Figure 4: ROC curves show that different peak detectors are better for different mark types. We show data for annotator TDH, and the same model ordering was observed across the other 3 annotators (Supplementary Figure S4). Note that even though we included all possible peaks (no threshold parameter $\lambda = 0$), some algorithms such as rseg and homer do not come close to 100% true positive rates. Also, some algorithms such as macs and macs.broad show non-monotonic ROC curves since their qvalue threshold parameter λ changes the size of each peak, in addition to the number of peaks.




Figure 5: Test error curves show mean and standard deviation over 100 randomly selected train and test sets, using the data sets with the most annotated windows for each mark type: H3K4me3_PGP Immune and H3K36me3_AM Immune (Table 3). Differences between models are clearly visible even in train data sets with just a few annotated windows. Similar model orderings were observed in other train/test splits (Supplementary Figure S5 and Supplementary Figure S6).

Table 2: Examples of different train and test sets. Note that the histone mark type is always the same (for example, we did not try to train a model on H3K4me3 and test it on H3K36me3).

Same	Train	Test
annotator, cell types	chr1	chr2
annotator, chromosomes	T cells	kidney cells
cell types, chromosomes	annotator TDH	annotator PGP

Table 3: Counts of windows, annotated regions, and samples of each cell type in each of the 7 data sets we analyzed.

mark	H3K36me3	H3K4me3	H3K36me3	H3K36me3	H3K4me3	H3K4me3	H3K4me3
annotator	TDH	TDH	AM	TDH	PGP	TDH	XJ
sample.set	other	other	immune	immune	immune	immune	immune
windows	4	29	23	4	30	27	12
noPeaks	72	536	752	230	1653	1656	702
peakStart	68	305	294	200	813	796	216
peakEnd	60	311	294	200	730	933	216
peaks		218	403		638	287	243
tcell			15	15	19	19	19
monocyte			5	5	6	6	6
bcell			1	1	2	2	2
kidney	1	1					
kidneyCancer	1	1					
skeletalMuscleCtrl	3	3					
skeletalMuscleMD	3	4					
leukemiaCD19CD10BCells		1					

The annotated region databases should be useful for benchmarking future peak detection algorithms. They can be viewed and downloaded at <http://cbio.ensmp.fr/~thocking/chip-seq-chunk-db/>

3.2 Train error for different mark types and annotators

We considered annotation data sets of two different histone mark types: H3K4me3 and H3K36me3. Is there any single algorithm that can fit data from both mark types? Or is each algorithm better for fitting one mark type or the other?

We first considered training each peak detection algorithm by choosing the threshold with minimum annotation error $\hat{\lambda}$ (3) on each data set (Supplementary Figure S1). The macs algorithm showed the minimum train error of 8.9–20% across the 4 annotated H3K4me3 data sets. The hmcan.broad algorithm showed the minimum train error of 7–20.2% across the 3 annotated H3K36me3 data sets.

We also computed ROC curves by varying the peak detection threshold λ for every algorithm (Figure 4). The ROC curves show that different algorithms are good detectors for different mark types. For example, it is clear that macs was the best peak detector for H3K4me3, and hmcan.broad was the best peak detector for H3K36me3. We observed similar trends in the ROC curves for the other annotation data sets (Supplementary Figure S2).

The ROC curves and the annotation error curves also show that the default parameter values $\tilde{\lambda}$ have generally higher false positive rates than the optimal parameter values $\hat{\lambda}$ chosen by minimizing the annotation error (3). For example, Supplementary Figure S3 shows two profiles where the default algorithms show false positives, but the trained algorithms achieve perfect peak detection.

3.3 Consistency of different annotators

When two experts annotate the same profiles, are they consistent? Does a peak detector trained on one annotator work well for another annotator? To answer these questions, we asked 4 different people to create annotated regions based on the same ChIP-seq profiles.

There were some windows in the H3K36me3 immune data sets that were annotated by both TDH and AM (Supplementary Figure S4). It is clear that these two annotators had consistent definitions of peak locations, but different levels of detail. For example, TDH used a peakStart and peakEnd region in many instances where AM used a peaks region. Also, there are other regions which only TDH annotated, and AM left un-annotated.

When we looked at train error of the different algorithms, we saw the same mark-specific model ordering that was independent of the annotator. For example, Supplementary Figure S2 shows similar ROC curves for the H3K4me3 immune annotation data sets, across 3 different annotators (TDH, PGP, and XJ). In particular, it is clear that macs is the best algorithm across the 3 annotators.

We also observed similar model orderings across annotators in terms of test error (Supplementary Figure S5). In particular, we trained models on annotations from TDH, and tested them on annotations from PGP (holding the mark H3K4me3 and immune cell types constant). We observed very little changes in test error when training on one or the other annotator.

Overall, these data indicate that annotated regions were consistent across annotators. Annotated regions can thus be used as a method for benchmarking peak detection algorithms, with results that are specific to the ChIP-seq data set, and quite robust across annotators.

3.4 Test on different regions of the same samples

Does a peak detector trained on just a few annotated chromosomes work well for detecting peaks on the other chromosomes of the same samples?

To answer this question, we performed cross-validation on each annotation data set by designating half of the windows as a test set. Then, we trained a model using only a few of the other windows, for training set sizes of from 1 to 12 windows. We randomly selected 100 train and test sets, trained each peak detector by picking the parameter $\hat{\lambda}$ with minimal error on the training set (3), then evaluated its peak detection by computing the annotation error on the test set of windows (2).

Figure 5 shows the test error as a function of training data size for two annotation data sets. We observed that each algorithm quickly achieves its model-specific minimum error, after only about 4 annotated windows in the train set. Furthermore, we observed that macs had the lowest test error for H3K4me3 data sets, and hmcan.broad had the lowest test error for H3K36me3 data sets.

Figure 5 also shows the test error of the default thresholds $\tilde{\lambda}$ suggested by the author of each algorithm. For some models such as sicer, the trained model $\hat{\lambda}$ is clearly better than the default model $\tilde{\lambda}$. For other models such as rseg, the trained model $\hat{\lambda}$ has about the same test error as the default model $\tilde{\lambda}$.

3.5 Test on different samples and cell types

Does a peak detector trained on some cell types work when applied to other cell types?

To answer this question, we trained each algorithm on several immune cell types (tcell, bcell, monocyte) and tested them on other cell types (kidney, kidney cancer, skeletal muscle, leukemia). We observed that most algorithms had similar test error when training on the same or different cell types (Supplementary Figure S6). The exceptions were the macs and macs.broad algorithms, which exhibited higher test error when training on samples of different cell types. These data suggest different samples have different kinds of background noise, and that it is important to annotate all samples of interest for accurate peak detection using the macs algorithm.

4 Discussion

4.1 Supervised versus unsupervised analysis

In the machine learning literature, a problem is considered supervised when there is a teacher or expert that provides correct predictions for a learning algorithm. In this paper, the type of supervision that we proposed was a database of annotated regions that represents where a scientist does and does not observe peaks. We used these annotated regions as a gold standard to define a supervised learning problem (2), which seeks a peak detector with minimal incorrect regions on a test data set. Furthermore, we proposed to train or calibrate standard ChIP-seq peak callers by choosing a threshold $\hat{\lambda}$ (3) that minimizes the number of incorrect regions.

In contrast, ChIP-seq peak detection without annotated regions can be considered an unsupervised learning problem. Usually, a default threshold $\hat{\lambda}$ and peak detection algorithm are first fit to a data set, and then the peaks are judged by visualizing them along with the data in a genome browser. If there are too many peaks the user can increase the threshold λ , or if there are not enough peaks the user can decrease the threshold λ .

In fact, our supervised analysis protocol with annotated regions is very similar, but is independent of any specific peak detection algorithm. First, one must visually inspect the ChIP-seq data in a genome browser, and create an annotated region database. After that, those annotations can be used to choose an appropriate model, and to calibrate its parameters.

4.2 Time required for annotation

Annotated regions are applicable to any ChIP-seq data analysis project, but their main weakness is the time it takes to create the annotation database. However, we found that it only takes about 10 minutes of manual visual inspection to find and annotate a whole window of several nearby regions.

Additionally, we were able to quickly create many regions by annotating dozens of samples at the same time. For example, the immune H3K4me3 sample set consists of 27 samples (Table 3). So when we found a region with a peak across all samples, we assigned the same peaks annotated region to all 27 samples.

As another example, it only took about 40 minutes to create the H3K36me3.TDH_other data set, which contains in total 8 samples and 200 annotated regions across 4 genomic windows (Table 3). And even though this was the smallest data set that we created, we were still able to observe clear differences in train error between the various algorithms (Supplementary Figure S1 and Supplementary Figure S2).

Also, the test error curves indicate that only a few annotated windows are necessary to calibrate peak detection thresholds λ (Figure 5). In both H3K4me3 and H3K36me3 data sets, the test error decreases to its model-specific minimum after annotating only about 4 windows.

Overall, these data indicate that a relatively small annotated region database of a few annotated windows across several samples is sufficient to calibrate and test peak detection algorithms.

5 Conclusions

We propose a supervised machine learning approach to ChIP-seq data analysis. Our approach involves first creating an annotated region database for a specific data set, and then using it to choose an appropriate peak detection algorithm.

We used this approach to benchmark the performance of several peak detectors on several H3K4me3 and H3K36me3 data sets. We observed that macs was the best peak detector for H3K4me3 data, and hmcen.broad was the best for H3K36me3 data. Furthermore, we observed these trends in train and test error, across several different annotators and cell types.

However, even the best peak detectors exhibited 10–20% test error (Figure 5), which could be improved. To develop even better peak detectors, we are interested in developing multi-parameter supervised learning algorithms, which have proven to be very successful in breakpoint detection [Rigaill et al., 2013].

We have made our annotated regions available as a public benchmark data set. Such annotated data sets are essential for making links between computational biology and the larger computer science research community. In particular, these annotated data sets make the ChIP-seq peak detection problem more accessible to machine learning researchers, who will now be able to work on developing supervised learning algorithms for peak detection. For example, new algorithms for broad marks should be tested alongside the hmcan.broad model, which we observed to be the best for H3K36me3 data. In the future, annotated region databases should be created to benchmark algorithms for other ChIP-seq experiments (e.g. H3K9me3).

We are interested in using systems such as Apollo [Lee et al., 2013], which extend genome browsers to support interactive annotation. We would also like to develop an interactive ChIP-seq data analysis system that displays optimal peak predictions which can be updated by adding annotated regions.

6 Acknowledgements

Thanks to David Bujold for help with the Epigenomes portal data and thanks to Haitham Ashoor, Chris Benner, Maxime Caron, Tao Liu, and Song Qiang for help with the peak callers.

6.1 Funding

This work was supported by computing resources provided by Calcul Quebec and Compute Canada, and by Canadian Institutes of Health Research grants EP1-120608 and EP1-120609, awarded to GB.

6.2 Conflict of interest statement.

None declared.

All of the sections after this page are to be considered supplementary materials.

A Source code and data links

A.1 Original ChIP-seq data profiles

<http://epigenomesportal.ca>

A.2 Annotated region database

<http://cbio.ensmp.fr/~thocking/chip-seq-chunk-db/>

A.3 R package with C code for computing the annotation error

<https://github.com/tdhock/PeakError>

A.4 \LaTeX and R source code for this article

<https://bitbucket.org/mugqic/chip-seq-paper>

B Example of annotated region data file

Below is an example of an annotation file that TDH created by manual visual inspection of the H3K4me3 immune cell samples (tcell, bcell, and monocyte). It is divided into 4 windows, each of which is separated by two returns/newlines. Each line represents an annotated region for several cell types. The cell types that are not listed all get a noPeaks annotation in the indicated region. Genomic regions were copied from the UCSC genome browser web page, and pasted into a text file.

```
chr11:118,092,641-118,095,026 peakStart monocyte tcell
chr11:118,095,334-118,096,640 peakEnd monocyte tcell
chr11:118,101,452-118,118,472 peaks
chr11:118,121,649-118,124,175 peaks monocyte tcell bcell

chr11:111,285,081-111,285,355 peakStart monocyte
chr11:111,285,387-111,285,628 peakEnd monocyte
chr11:111,299,681-111,337,593 peaks
chr11:111,635,157-111,636,484 peakStart monocyte tcell bcell
chr11:111,637,473-111,638,581 peakEnd monocyte tcell bcell

chr1:32,717,194-32,721,976 peaks tcell
chr1:32,750,608-32,756,699 peaks
chr1:32,757,261-32,758,801 peaks tcell bcell monocyte

chr2:26,567,544-26,568,406 peakStart bcell tcell monocyte
chr2:26,568,616-26,568,862 peakEnd bcell tcell monocyte
chr2:26,569,573-26,571,905 peakEnd bcell tcell monocyte
chr2:26,578,595-26,632,223 peaks
chr2:26,634,282-26,636,118 peaks monocyte
```

C Definition of the annotation error

In this section we give the precise mathematical definition of the annotation error.

C.1 Data definition

Let there be n annotated training samples, all of the same histone mark type. For simplicity, and without loss of generality, let us consider just one chromosome with d base pairs. Let $\mathbf{x}_1 \in \mathbb{R}^d, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be the vectors of coverage across that chromosome. For example $d = 249,250,621$ is the number of base pairs on chr1, and $\mathbf{x}_i \in \mathbb{R}^d$ is the H3K4me3 coverage profile on chr1 for one sample $i \in \{1, \dots, n\}$.

We also have exactly 4 sets of annotated regions $\underline{R}_i, \overline{R}_i, R_i^+, R_i^-$ for each sample $i \in \{1, \dots, n\}$:

Data	Type	Color	Symbols
Coverage			$\mathbf{x}_1, \dots, \mathbf{x}_n$
Weak annotations	peaks	■	R_1^+, \dots, R_n^+
	noPeaks	■	R_1^-, \dots, R_n^-
Strong annotations	peakStart	■	$\underline{R}_1, \dots, \underline{R}_n$
	peakEnd	■	$\overline{R}_1, \dots, \overline{R}_n$

For each sample $i \in \{1, \dots, n\}$, R_i^+ is a set of regions, and each region is an interval of base pairs. For example, Figure 2 shows a sample with one annotation of each type.

A peak detection function or peak caller $c : \mathbb{R}^d \rightarrow \{0, 1\}^d$ gives a binary peak call prediction $\mathbf{y} = c(\mathbf{x}) \in \{0, 1\}^d$ given some coverage profile \mathbf{x} .

The goal is to learn how to call peaks $c(\mathbf{x}_i)$ which agree with the annotated regions $R_i^+, R_i^-, \underline{R}_i, \overline{R}_i$ for some test samples i . To quantify the error of the peak calls with respect to the annotation data, we define the following functions.

C.2 Weak annotation error

The weak annotations R_i^+, R_i^- count whether there are any peaks overlapping a region. They are called weak because each peaks region can only produce a false negative (but not a false positive), and each noPeaks region can only produce a false positive (but not a false negative). Let

$$B(\mathbf{y}, \mathbf{r}) = \sum_{j \in \mathbf{r}} y_j \quad (4)$$

be the number of bases which have peaks overlapping region \mathbf{r} . Then for a sample i , the number of weak false positives is

$$\text{WFP}(\mathbf{y}, R_i^-) = \sum_{\mathbf{r} \in R_i^-} I[B(\mathbf{y}, \mathbf{r}) > 0] \quad (5)$$

and the number of weak true positives is

$$\text{WTP}(\mathbf{y}, R_i^+) = \sum_{\mathbf{r} \in R_i^+} I[B(\mathbf{y}, \mathbf{r}) > 0], \quad (6)$$

where I is the indicator function.

C.3 Strong annotation error

The strong annotations $\underline{R}_i, \overline{R}_i$ count the number of peak starts and ends occurring in the given regions. They are called strong because each region can produce a false positive (more than 1 peak start/end predicted in

the region) or a false negative (no peak start/end predicted). First, let $y_0 = y_{d+1} = 0$ and define the set of first bases of all peaks as

$$\underline{\mathcal{I}}(\mathbf{y}) = \{j \in \{1, \dots, d\} : y_j = 1 \text{ and } y_{j-1} = 0\} \quad (7)$$

and the set of last bases of all peaks as

$$\overline{\mathcal{I}}(\mathbf{y}) = \{j \in \{1, \dots, d\} : y_j = 1 \text{ and } y_{j+1} = 0\}. \quad (8)$$

For a sample i , the number of strong false positives is

$$\text{SFP}(\mathbf{y}, \underline{R}_i, \overline{R}_i) = \sum_{\mathbf{r} \in \underline{R}_i} I[|\mathbf{r} \cap \underline{\mathcal{I}}(\mathbf{y})| > 1] + \sum_{\mathbf{r} \in \overline{R}_i} I[|\mathbf{r} \cap \overline{\mathcal{I}}(\mathbf{y})| > 1], \quad (9)$$

and the number of strong true positives is

$$\text{STP}(\mathbf{y}, \underline{R}_i, \overline{R}_i) = \sum_{\mathbf{r} \in \underline{R}_i} I[|\mathbf{r} \cap \underline{\mathcal{I}}(\mathbf{y})| > 0] + \sum_{\mathbf{r} \in \overline{R}_i} I[|\mathbf{r} \cap \overline{\mathcal{I}}(\mathbf{y})| > 0]. \quad (10)$$

C.4 Total annotation error

For a sample i , the total number of false positives is

$$\text{FP}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^-) = \text{WFP}(\mathbf{y}, R_i^-) + \text{SFP}(\mathbf{y}, \underline{R}_i, \overline{R}_i), \quad (11)$$

the total number of true positives is

$$\text{TP}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+) = \text{WTP}(\mathbf{y}, R_i^+) + \text{STP}(\mathbf{y}, \underline{R}_i, \overline{R}_i), \quad (12)$$

the total number of false negatives is

$$\text{FN}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+) = |\underline{R}_i| + |\overline{R}_i| + |R_i^+| - \text{TP}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+). \quad (13)$$

The annotation error E quantifies the number of incorrect annotated regions:

$$E(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+, R_i^-) = \text{FP}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^-) + \text{FN}(\mathbf{y}, \underline{R}_i, \overline{R}_i, R_i^+). \quad (14)$$

The annotation error E can be easily computed using the C code in the R package PeakError on GitHub: <https://github.com/tdhock/PeakError>.

C.5 ROC analysis

A standard unsupervised peak caller can be characterized as a function $c_\lambda : \mathbb{R}^d \rightarrow \{0, 1\}^d$, where the significance threshold $\lambda \in \mathbb{R}$ controls the number of peaks detected. Receiver Operating Characteristic (ROC) curves can also be used to show the train error of all thresholds of a peak detector on an annotation data set. Define the false positive rate as

$$\text{FPR}(\lambda) = \frac{\sum_{i=1}^n \text{FP}[c_\lambda(\mathbf{x}_i), \underline{R}_i, \overline{R}_i, R_i^-]}{|\underline{R}_i| + |\overline{R}_i| + |R_i^-|} \quad (15)$$

and the true positive rate as

$$\text{TPR}(\lambda) = \frac{\sum_{i=1}^n \text{TP}[c_\lambda(\mathbf{x}_i), \underline{R}_i, \overline{R}_i, R_i^+]}{|\underline{R}_i| + |\overline{R}_i| + |R_i^+|}. \quad (16)$$

ROC curves are traced by plotting $\text{TPR}(\lambda)$ versus $\text{FPR}(\lambda)$ for all possible values of the peak detection threshold λ .

D Details of peak calling algorithms

Zhang et al. [2008] proposed Model-based Analysis of ChIP-Seq (MACS). We downloaded MACS version 2.0.10 12162013 (commit ca806538118a85ec338674627f0ac53ea17877d9 on GitHub). The significance threshold λ is the qvalue cutoff parameter `-q`. We used a grid of 59 qvalue cutoffs from 0.8 to 10^{-15} .

The macs broad peak caller is more difficult to train because it has 2 different qvalue cutoff parameters (`-q` and `--broad-cutoff`). So we used the same grid of 59 qvalue cutoffs for `-q`, and then defined

$$\text{broad-cutoff} = q \times 1.122. \quad (17)$$

The default macs algorithms use a default qvalue cutoff of $\tilde{\lambda} = 0.05$.

Ashoor et al. [2013] proposed the HMCAN algorithm which uses GC-content and copy number normalization. We downloaded HMCAN commit 9d0a330d0a873a32b9c4fa72c94d00968132b9ef from BitBucket. We used the default GC content normalization file provided by the authors. We used two different parameter files to test two different peak detectors:

name	mergeDistance
hmcan	200
hmcan.broad	1000

We then ran HMCAN with `finalThreshold=0`, and defined λ as a threshold on column 5 in the regions.bed file. Both hmcan and hmcan.broad use a default `finalThreshold` of $\tilde{\lambda} = 10$.


Song and Smith [2011] proposed the RSEG algorithm. We downloaded RSEG version 0.4.8 from <http://smithlabresearch.org/software/rseg/>. Upon recommendation of the authors, we saved computation time by running `rseg-diff` using options `-training-size 100000` and `-i 20`. We used the `-d` option to specify a dead regions file for hg19 based on our alignment pipeline. We defined the significance threshold λ as the sum of posterior scores (column 6 in the output .bed file). For the default RSEG algorithm, we used all the peaks in the output file, meaning a posterior score threshold of $\tilde{\lambda} = 0$.

Zang et al. [2009] proposed the SICER algorithm. We downloaded SICER version 1.1 from http://home.gwu.edu/~wpeng/SICER_V1.1.tgz. We defined the significance threshold λ as the FDR (column 8) in the islands-summary output file. For the default SICER algorithm, we used an FDR of $\tilde{\lambda} = 0.01$ as suggested in the README and example.

Heinz et al. [2010] proposed the HOMER set of tools for DNA motif detection. We used the `findPeaks` program in HOMER version 4.1 with the `-style histone` option. We defined the significance threshold λ as the “p-value vs Control” column. We defined the default model as all peaks in the output file.


E Supplementary Figures

Supplementary Figure S1: annotation error curves for all algorithms and data sets

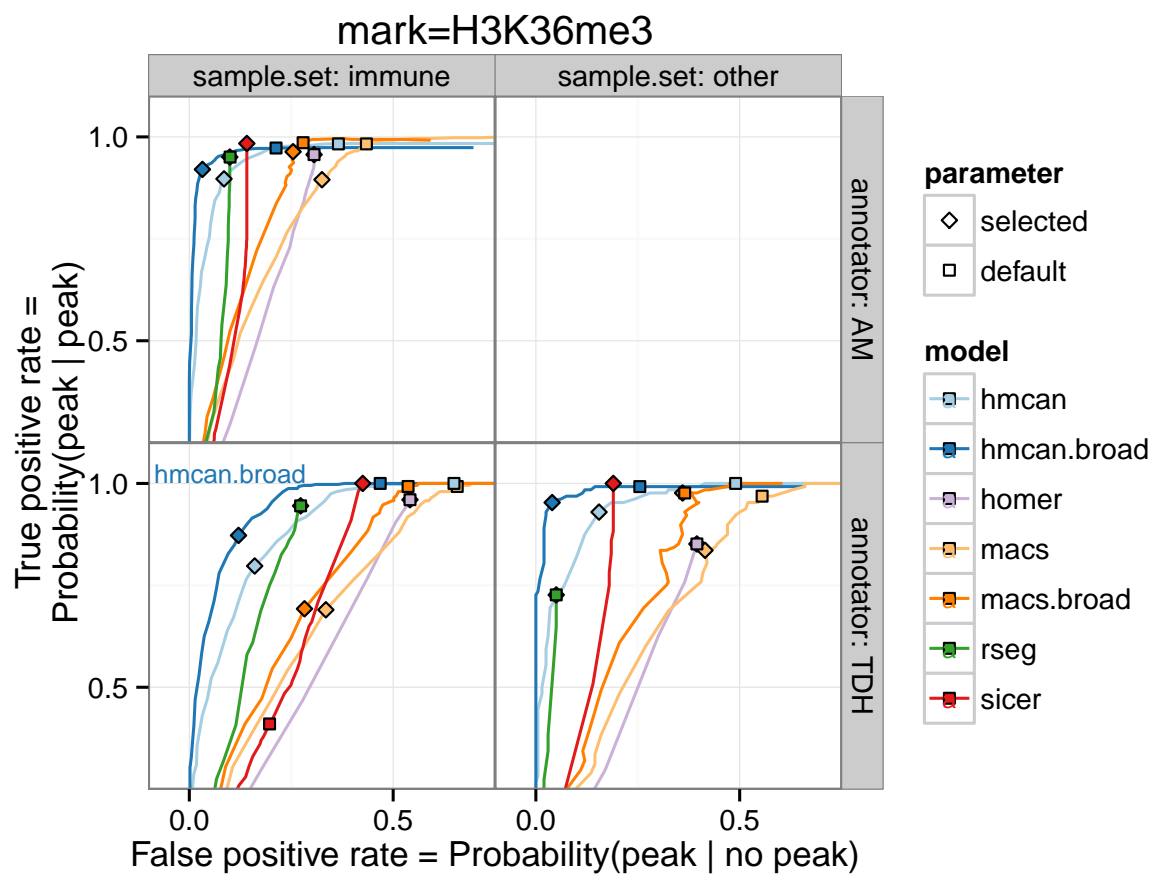


Train error curves for each data set (rows) and each algorithm (columns). A vertical red line shows the default parameter for each model, and a horizontal grey line shows the minimum error model for each data set.

Supplementary Figure S2: ROC curves for train error of all algorithms on all data sets

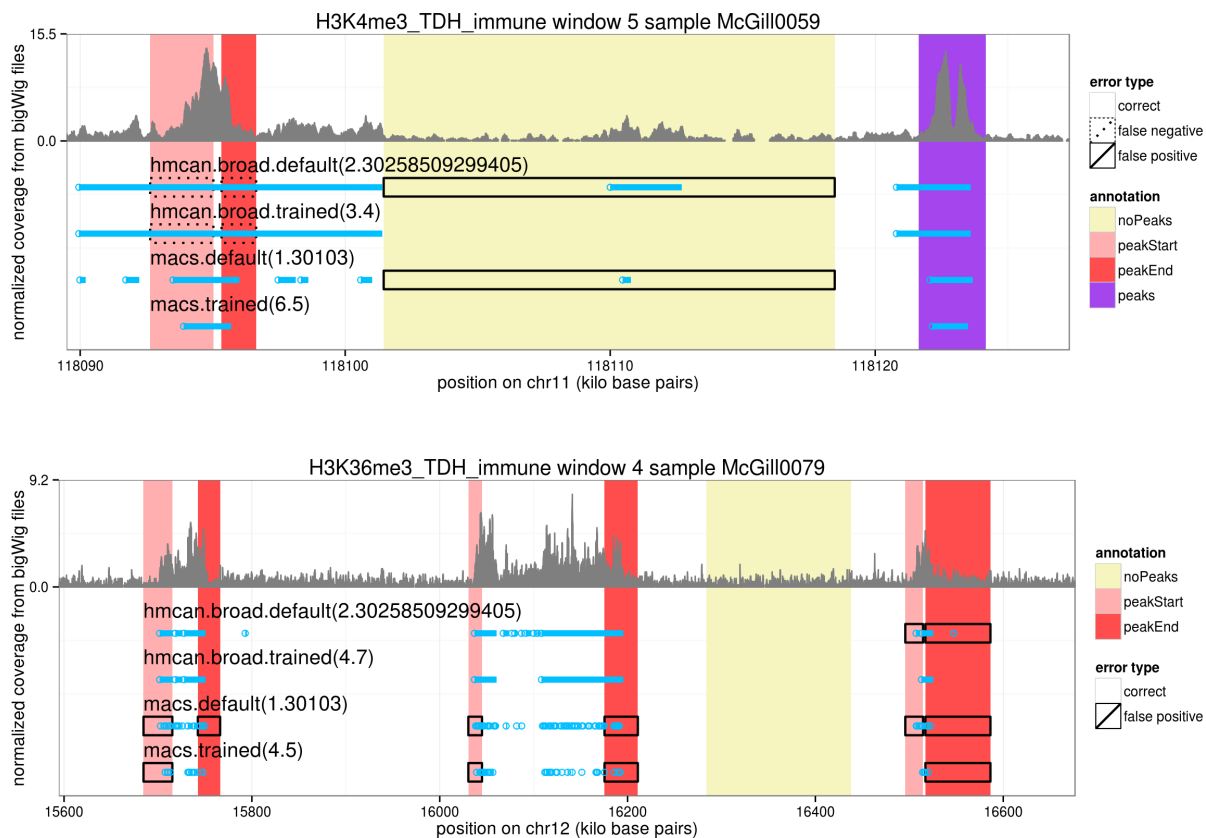


It is clear that macs is the best for H3K4me3 data.



It is clear that hmcen.broad is the best for H3K36me3 data.

Supplementary Figure S3: examples of fitted peak detectors




Shown are default and trained algorithms for samples of 2 different histone mark types. Consistent with the quantitative results, it is clear that macs.trained is the best for the H3K4me3 data, and that hmcan.broad.trained is the best for the H3K36me3 data.


Supplementary Figure S4: Windows labeled by 2 different annotators

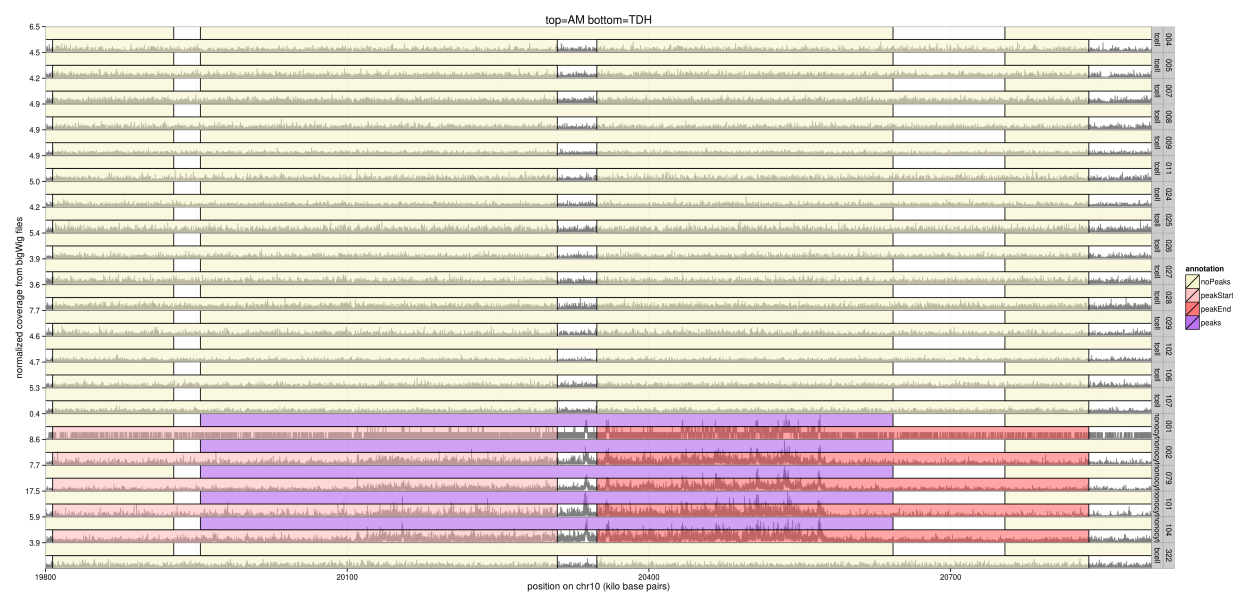
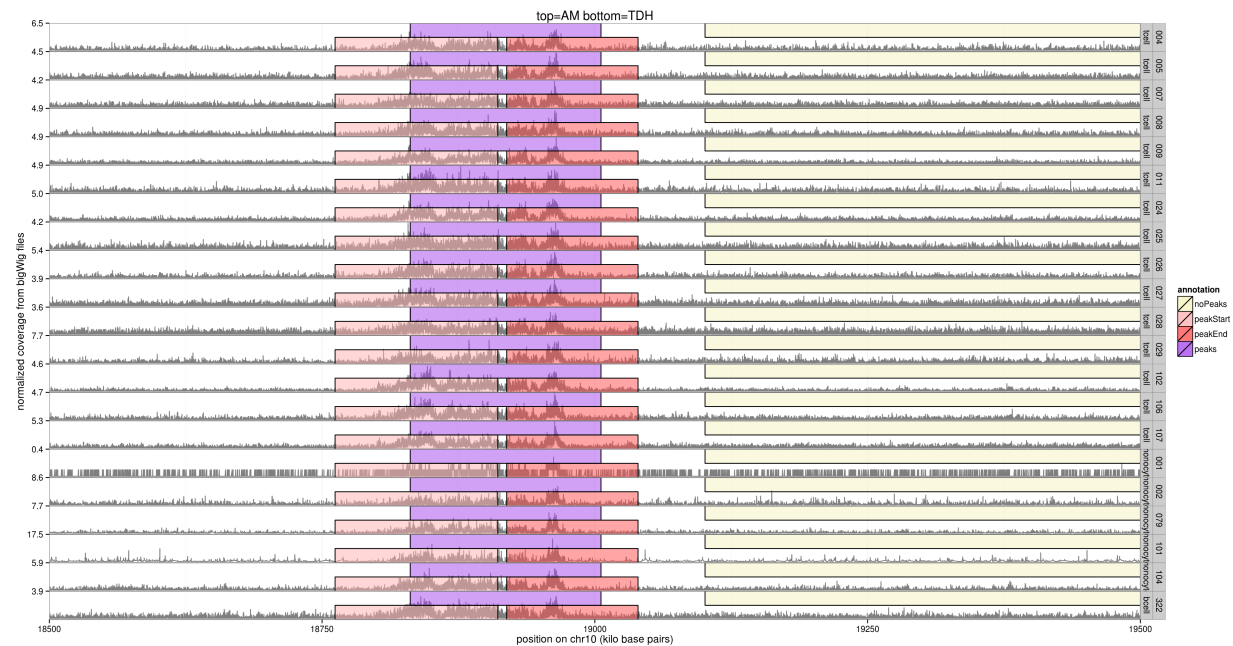
Annotated regions for 2 different annotators (top=AM, bottom=TDH) on the same genomic regions and sample sets (H3K36me3 immune cell types). It is clear that the annotators focus on different levels of detail, but have in general the same visual definition of a peak.



First, we show just 1 profile:




We also show all profiles, for several annotated windows:








Supplementary Figure S5: train on one annotator, test on another annotator



H3K4me3 immune cell data sets were annotated by TDH and PGP. We trained models on one or the other annotator, and then tested those models on the same or a different annotator. We repeated 100 trials of the following computational experiment: we randomly designated half of the annotated windows as test data in each data set, then selected the parameter $\hat{\lambda}$ with minimum annotation error on a train set of 10 annotated windows.

It is clear for all models that it does not make much difference in terms of test error when training on one or another annotator.

Supplementary Figure S6: train on some cell types, test on some other cell types



TDH annotated H3K4me3 data sets of immune and other cell types. We trained models on one or the other cell types, and then tested those models on the same or a different set of cell types. We repeated 100 trials of the following computational experiment: we randomly designated half of the annotated windows as test data in each data set, then selected the parameter $\hat{\lambda}$ with minimum annotation error on a train set of 10 annotated windows.

It is clear that for some models such as *macs* and *macs.broad*, a model trained on the different cell types (e.g. right panel: train on immune, test on other) yields higher test error than a model trained on the same cell types (e.g. train on other, test on other). It is also clear that the train data do not make much difference for other models, such as *rseg*, *sicer*, *hmcan*, *hmcan.broad*, and *homer*.

References

- H. Ashoor, A. Hrault, A. Kamoun, F. Radvanyi, V. B. Bajic, E. Barillot, and V. Boeva. HMCAN: a method for detecting chromatin modifications in cancer samples using ChIP-seq data. *Bioinformatics*, 29(23): 2979–2986, 2013. doi: 10.1093/bioinformatics/btt524.
- T. Bailey, P. Krajewski, I. Ladunga, C. Lefebvre, Q. Li, T. Liu, P. Madrigal, C. Taslim, and J. Zhang. Practical guidelines for the comprehensive analysis of ChIP-seq data. *PLoS computational biology*, 9(11): e1003326, 2013.
- A. Barski, S. Cuddapah, K. Cui, T.-Y. Roh, D. E. Schones, Z. Wang, G. Wei, I. Chepelev, and K. Zhao. High-resolution profiling of histone methylations in the human genome. *Cell*, 129(4):823–837, 2007.
- K.-B. Chen and Y. Zhang. A varying threshold method for chip peak-calling using multiple sources of information. *Bioinformatics*, 26(18):i504–i510, 2010. doi: 10.1093/bioinformatics/btq379.
- D. U. Gorkin, D. Lee, X. Reed, C. Fletez-Brant, S. L. Bessling, S. K. Loftus, M. A. Beer, W. J. Pavan, and A. S. McCallion. Integration of chip-seq and machine learning reveals enhancers and a predictive regulatory sequence vocabulary in melanocytes. *Genome research*, 22(11):2290–2301, 2012.
- S. Heinz, C. Benner, N. Spann, E. Bertolino, Y. C. Lin, P. Laslo, J. X. Cheng, C. Murre, H. Singh, and C. K. Glass. Simple Combinations of Lineage-Determining Transcription Factors Prime cis-Regulatory Elements Required for Macrophage and B Cell Identities. *Molecular cell*, 38(4):576–589, 2010.
- T. D. Hocking, G. Schleiermacher, I. Janoueix-Lerosey, V. Boeva, J. Cappel, O. Delattre, F. Bach, and J.-P. Vert. Learning smoothing models of copy number profiles using breakpoint annotations. *BMC Bioinformatics*, 14(164), May 2013.
- T. D. Hocking, V. Boeva, G. Rigau, G. Schleiermacher, I. Janoueix-Lerosey, O. Delattre, W. Richer, F. Bourdeaut, M. Suguro, M. Seto, F. Bach, and J.-P. Vert. SegAnnDB: interactive Web-based genomic segmentation. *Bioinformatics*, 2014. doi: 10.1093/bioinformatics/btu072.
- T. R. Jones, A. E. Carpenter, M. R. Lamprecht, J. Moffat, S. J. Silver, J. K. Grenier, A. B. Castoreno, U. S. Eggert, D. E. Root, P. Golland, et al. Scoring diverse cellular morphologies in image-based screens with iterative feedback and machine learning. *Proceedings of the National Academy of Sciences*, 106(6): 1826–1831, 2009.
- W. Kent, C. Sugnet, T. Furey, K. Roskin, T. Pringle, A. Zahler, and D. Haussler. The human genome browser at UCSC. *Genome Research*, 12(6):996–1006, June 2002.
- W. J. Kent, A. S. Zweig, G. Barber, A. S. Hinrichs, and D. Karolchik. Bigwig and bigbed: enabling browsing of large distributed datasets. *Bioinformatics*, 26(17):2204–2207, 2010. doi: 10.1093/bioinformatics/btq351.
- E. Lee, G. A. Helt, J. T. Reese, M. C. Munoz-Torres, C. P. Childers, R. M. Buels, L. Stein, I. H. Holmes, C. G. Elisk, and S. E. Lewis. Web Apollo: a web-based genomic annotation editing platform. *Genome Biol*, 14:R93, 2013.
- C. B. Nielsen, H. Younesy, H. O’Geen, X. Xu, A. R. Jackson, A. Milosavljevic, T. Wang, J. F. Costello, M. Hirst, P. J. Farnham, et al. Spark: a navigational paradigm for genomic data exploration. *Genome research*, 22(11):2262–2269, 2012.
- H. U. Osmanbeyoglu, R. J. Hartmaier, S. Oesterreich, and X. Lu. Improving chip-seq peak-calling for functional co-regulator binding by integrating multiple sources of biological information. *BMC genomics*, 13(Suppl 1):S1, 2012.

- G. Rigai, T. D. Hocking, F. Bach, and J.-P. Vert. Learning sparse penalties for change-point detection using max margin interval regression. In S. Dasgupta and D. McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, ICML '13, New York, NY, USA, June 2013. ACM.
- Q. Song and A. D. Smith. Identifying dispersed epigenomic domains from chip-seq data. *Bioinformatics*, 27(6):870–871, 2011. doi: 10.1093/bioinformatics/btr030.
- A. M. Szalkowski and C. D. Schmid. Rapid innovation in ChIP-seq peak-calling algorithms is outdistancing benchmarking efforts. *Briefings in Bioinformatics*, 12(6):626–633, 2011. doi: 10.1093/bib/bbq068.
- H. Younesy, C. Nielsen, T. Mller, O. Alder, R. Cullum, M. Lorincz, M. Karimi, and S. Jones. An interactive analysis and exploration tool for epigenomic data. *Computer Graphics Forum*, 32(3pt1):91–100, 2013. ISSN 1467-8659. doi: 10.1111/cgf.12096.
- C. Zang, D. E. Schones, C. Zeng, K. Cui, K. Zhao, and W. Peng. A clustering approach for identification of enriched domains from histone modification ChIP-Seq data. *Bioinformatics*, 25(15):1952–1958, 2009. doi: 10.1093/bioinformatics/btp340.
- Y. Zhang, T. Liu, C. A. Meyer, J. Eeckhoutte, D. S. Johnson, B. E. Bernstein, C. Nusbaum, R. M. Myers, M. Brown, W. Li, et al. Model-based analysis of ChIP-Seq (MACS). *Genome Biol*, 9(9):R137, 2008.